

DT05 Rec'd PCT/PTO 30 NOV 2004

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE IPEA/US

International Appln. No.: PCT/US02/17099
International Filing Date: 31 May 2002
Priority Date Claimed: 31 May 2002
Title: IMPROVING TRANSACTION-
PROCESSING PERFORMANCE
BY PREFERENTIALLY REUSING
FREQUENTLY USED PROCESSES
Applicant's Reference: PCT 9/120,000 (Dkt. 544-L)
Applicant: Unisys Corporation

May 7, 2004

RESPONSE TO WRITTEN OPINION

Mail Stop PCT, ATTN: IPEA/US Via Facsimile No. 703 305-3230
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

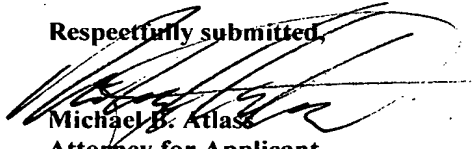
Attention: Jack B. Harvey
Authorized Officer

Dear Sir:

In response to the Written Opinion dated 12 March 2004, issued in connection with the above-captioned International Application, Applicants' response is attached hereto.

Attached is a copy of page 16 of the specification which indicates claim 4 as being deleted. Claim 4 is deleted since it is essentially the same as claim 2.

Respectfully submitted,



Michael B. Atlase
Attorney for Applicant
Reg. No. 30,606

Unisys Corporation
Unisys Way, E8 114
Blue Bell, PA 19424-001 US
Telephone: (215) 986-4111
Facsimile: (215) 986-3090

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for scheduling transactions within a computing system, comprising the steps of:

- 5 - receiving a transaction request from a client,
 - determining the idle server process that most recently finished execution of a previous transaction, and forwarding the transaction request to the most recently idle process.

10 2. A method in accordance with claim 1, further comprising the step of storing the most recent finish time and the status of a server process.

3. A method in accordance with claim 1 or 2, whereby the method is implemented as a sub-module within a gateway process module, whereby the gateway module is responsible for receiving transaction requests from a client, and allocating the requests to server processes.

15 4. A method in accordance with claim 3, the method further comprising the step of storing the most recent finish time and the status of a server process.

DELETE

20 5. A method in accordance with any one of the preceding claims, wherein the method is implemented as a software module arranged to operate in an enterprise application environment.

25 6. A method in accordance with any one of the preceding claims, wherein each server process interacts with a separate central processing unit.

7. A system for scheduling transactions within a computing system, comprising:

- 30 - a transaction request means arranged to receive a transaction request from a client;

MATERIAL FOR RESPONSE TO PCT
REGARDING INTERNATIONAL SEARCH REPORT
DATED: 31 MAY 2002 (31.05.2002)

APPLICANTS' FILE REFERENCE:	PCT 09/120,000
INTERNATIONAL APPLICATION NO:	PCT/US02/17099
APPLICANT:	UNISYS CORPORATION

DOCUMENTS CONSIDERED TO BE RELEVANT

These are category Y: To U.S. Patent 6,185,619B1 (Joffe, et al) 06 February 2001 -- relevant to claims 1-9.

Category Y: U.S. 5,918,017A (Attanasio, et al) 29 June 1999 -- relevant to claims 1-9.

International Filing Date: 31 May 2002.

Date of Actual Completion of International Search: 12 September 2002.

In response to the Examiner's International Search Report and Written Opinion on the patent application entitled: "IMPROVING TRANSACTION-PROCESSING PERFORMANCE BY PREFERENTIALLY RE-USING FREQUENTLY USED PROCESSES", the following response is made to the Examiner's Written Opinion.

1. On review of the claims, it has been deemed that claim 4 is redundant, as it claims a method step which is substantially identical to claim 2. As a result, the attached page of claims will indicate that claim 4 has been crossed over and marked "delete", thus claim 4 need no longer be considered as part of this application.

2. In the Written Opinion, the Examiner had contended that claims 1-9 are not inventive in light of Joffe, et al. (U.S. 6,185,619), in view of Attanasio, et al. (U.S. 5,918,017).

A review of these cited documents has indicated that these documents do not disclose the essential features of the claimed invention, nor do they even teach directionally toward the claimed invention. Rather, Applicants' would assert that both of the cited documents actually teach away from the claimed invention.

3. Both of these prior art documents to Joffe and Attanasio, are used to broadly describe systems that supposedly efficiently allocate transactions in a multi-computer system.

Joffe describes a system which is designed to route requests for data objects (for example, a request from a user to view of Web page) over to a "best server" based on a pre-determined routing policy, as will be generally seen in Joffe, column 3, lines 44-60. This pre-determined routing policy may take into account one of a number of factors, including routing the request to a server that has:

- (a) the least number of open TCP connections;
- (b) the most available free RAM;
- (c) the most available free SWAP (Virtual Memory);
- (d) the highest amount of CPU idle time; or
- (e) the fastest ICMP route to the client's machine.

(Here, notice column 2, line 64 to column 3, line 3).

This pre-determined routing policy is based on simple common sense and on knowledge which would be in the possession of a skilled addressee. Thus namely, the principle that a server with the least load or the longest idle time is generally the server to which the next queued transaction should be sent. This is the standard conventional wisdom, and most routing policies are based on such logic.

4. Now, as to Attanasio, here quite similarly, Attanasio describes a process for dynamically altering the load on a cluster of computers by measuring the current server load metrics according to a "configurable policy". (Here, notice column 2, lines 3-9).

In Attanasio at column 7, lines 19-62, there are examples of possible metrics provided, which can include the level of utilization of the memory buffers, and forward delay times, etc. These raw metric values are then provided to a "weights" function and the Attanasio invention "adjusts the weights of the Executor 610 routing algorithm" so that servers which are "lightly loaded" according to the load metrics, will receive a larger portion of the incoming TCP connections (See column 8, lines 22-25 of Attanasio).

Therefore, the underlying principle of operation of the invention of Attanasio is quite identical to the underlying principle in Joffe — namely, the server with the lowest load, or the longest idle time, is the server that receives the next queued transaction requests.

5. Taking an overall look and noticing that in other words, the longer a server has been idle or the "less" it has been used in a given time period, the more likely that the next job will be routed to that particular server.

6. Now, very much in contrast, the present invention relies on a principle that is anti-intuitive, and quite surprising, and which produces an unexpected benefit.

The Applicants have surprising found that the server that was "most recently made idle", that is to say, the server with the lowest idle time — is the most appropriate server to receive the next queued transaction request.

As stated in Applicants' specification, this is due to the fact that high-end transaction servers have a number of peculiar characteristics:

Firstly, with the advent of faster processors, most of the time spent in processing a transaction occurs in loading/unloading registers and cache memory, rather than spending time in the actual processor "computation".

Secondly, the transaction processing systems in many cases, repeat similar transactions in a given time period.

Therefore, much of the data which is already held in registers and in cache memory (which is normally over-written) could be re-

used, saving the total amount of time required to process a transaction.

This feature is clearly spelled-out in each of the independent claims which state that the transaction request is forwarded to -- "the most recently idle process" --. For an example here, see the last line of claims 1, 7, 8 and 9.

7. Either of the cited references to Joffe or to Attanasio, either individually or taken together, will be seen not to teach or even suggest a methodology for assigning transaction requests to one of a plurality of servers or processors in accordance with the presently claimed invention.

Both documents cited as references, and indeed the common general knowledge in the art, do not teach such a methodology, just because it runs counter to the established industry practice and to "common sense".

Since there is no teaching in either the cited prior art or even in the common general knowledge that such a methodology is useful, Applicants would submit that the claims as stated are inventive in spite of the references to Joffe and Attanasio.

8. Therefore, in this regard, Examiner is encouraged to take a further look at the presently existing claims of Applicants, and differentiate the functionality that is provided by Applicants which cannot be accomplished or taught by the cited references. Thus, it should be concluded that the extant claims do provide an inventive configuration.